Contribution ID: **33**   Type: **not specified**

The development of hybrid metadata storage for PanDA Workload Management System

Friday, 2 October 2015 15:55 (15 minutes)

Scientific computing in a field of High Energy and Nuclear Physics (HENP) produces vast volumes of data. ATLAS, one of the largest collaborations ever assembled in the sciences, is at the forefront of research at Large Hadron Collider (LHC), operating at the international CERN Laboratory in Geneva, Switzerland, daily runs up to 1.5 M jobs and submit them using PanDA workload management system. For tracking the execution of computational and analytical tasks PanDA uses monitor application, which contains a set of summary tables, charts and graphs, aggregating data from the central SQL-based metadata storage (RDBMS Oracle). The growth rate of the volume of stored information has increased significantly over the last few years: from 500 hundreds completed jobs per day in 2011 up to 2 million during LHC Run 1 (2012-2013). Present metadata storage technology significantly limits the analytical tasks performance.

This research work is focused on the development of a Hybrid Metadata Storage Framework (HMSF) that would improve scalability and performance of PanDA metadata store. In this framework, the scalability issue is addressed by integrating relational database and NoSQL data store, which combines the strengths of both. We have developed a prototype of HMSF that provides data transfer and synchronization between parts of hybrid storage, with Cassandra as NoSQL backend. HMSF have an API providing interface, which interprets requests from external applications. PanDA monitor was partly adopted to interact with HMSF. The operational data queries are forwarded to the primary SQL-based repository and the analytic data requests are processed by NoSQL database, which stores prepared query-specific data structures.

The performance and scalability tests of HMSF-adopted part of PanDA monitor shows that data aggregation and precalculation in advance, with the help of HMSF synchronization mechanisms, provide significant performance improvement without adding much complexity to the resulting system.

Primary author: Ms GRIGORIEVA, Maria (National Research Center "Kurchatov Institute")

Co-authors: Dr KLIMENTOV, Alexei (Brookhaven National Lab); Prof. KAUSHIK, De (University of Texas at Arlington); Mr RYABINKIN, Eygene (National Research Center "Kurchatov Institute", Moscow, Russia); Ms GOLOSOVA, Marina (National Research Center "Kurchatov Institute", Moscow, Russia)

Presenter: Ms GRIGORIEVA, Maria (National Research Center "Kurchatov Institute")

Session Classification: Workload Management Systems in Applied Research and BigData